

IN THE CLAIMS

1. (Currently Amended) A method for scheduling prefetches into a Least Recently Used (LRU) cache of a data storage system, the method comprising:

remotely modeling dynamic operation of the LRU cache in a model, the model including a model of data elements currently stored within the LRU cache;

assigning a priority value to modeled data elements according to their history;

assigning a priority value to a requested data element based at least partially on whether a preceding data element is present in the LRU cache;

making a cache management decision based upon the model comprising;

intercepting a request for a data element from a stream of Input/Output (I/O) data requests sent from a host and addressed to the LRU cache; and

determining whether to schedule a prefetch of a data element logically successive to the requested data element in accordance with contents of the LRU cache as indicated by the remote model; and

executing prefetches into the LRU cache in response to select cache management decisions.

2. (Canceled).

3. (Canceled).

4. (Currently Amended) The method of claim ~~2~~ 1, wherein the LRU cache is a native LRU-only cache, and further comprising the step of leaving the native LRU-only cache substantially unmodified while conducting the steps of claim ~~2~~ 1.

5- 6. (Canceled).

7. (Currently Amended) The method of claim 1, wherein remotely modeling the LRU cache further comprises:

determining a size of the LRU cache;

periodically fetching an LRU cache I/O rate of from the LRU cache; and

periodically fetching an LRU cache hit rate of from the LRU cache.

8. (Original) The method of claim 1, wherein remotely modeling the cache further comprises periodically calculating a single reference residency time (SRRT) for a data element within the cache.

9-11. (Canceled).

12. (Previously Presented) The method of claim 1, wherein assigning a priority value further comprises assigning a priority value comprising the priority value assigned to the preceding data element plus one when the preceding data element is found to be present in the cache.

13. (Previously Presented) The method of claim 1, wherein determining whether to schedule a prefetch of a data element further comprises comparing the priority value of the requested element with a dynamic threshold.

14. (Original) The method of claim 13, further comprising prefetching the requested data element into the cache if the priority value of the requested data element is greater than the dynamic threshold.

15. (Previously Presented) The prefetch method of claim 1, further comprising periodically reevaluating the performance of the remote model.

16. (Currently Amended) The method of claim 15, wherein periodically reevaluating the performance of the ~~cache~~ remote model further comprises determining if ~~the~~ a dynamic threshold used in the remote model of the cache accurately models the performance of the cache.

17. (Original) The method of claim 16, wherein determining if the dynamic threshold accurately models the performance of the cache comprises comparing the performance of the dynamic threshold with an alternate dynamic threshold.

18. (Currently Amended) The method of claim 15, further comprising automatically updating ~~the~~ a dynamic threshold used in the remote model of the cache when another dynamic threshold is deemed to be more effective.

19. (Original) The method of claim 1, wherein making a cache management decision comprises deciding to schedule a prefetch, and further comprising scheduling a prefetch by sending an I/O request to the cache.

20. (Original) A method for scheduling prefetches in a data storage system having a host and a cache, the method comprising the steps of:

providing a cache for caching Input/Output (I/O) data;

providing a prefetch module remote to the cache;

remotely modeling the cache within the prefetch module and determining whether to schedule a prefetch of data into the cache according to the results of the step of remotely modeling the cache, the step of remotely modeling the cache module further comprising:

examining the history of a data element in the cache;

assigning a priority value to the data element according to its history;

comparing that priority value to a predetermined threshold value;

determining a size of memory used in the cache;

periodically fetching an I/O rate of the cache from the cache;

periodically fetching a hit rate of the cache from the cache; and

determining a single reference residency time for a data element within the cache;

intercepting a stream of I/O information from the host to the cache to locate a requested data element;

determining if the requested data element in the stream of I/O information is already present within the cache;

making the requested data element a youngest member of the cache;

determining if the data element preceding the requested data element is present in the cache;

assigning a priority value to the requested data element;

periodically reevaluating the performance of the cache versus an internal model of the cache if the number of I/O requests received by the cache is greater than a predetermined number;

updating the dynamic threshold used in the internal model of the cache if the dynamic threshold value does not adequately model the performance of the cache;

comparing the priority value of the requested data element with the dynamic threshold value;

and

prefetching the requested data element if the priority value of the requested data element is greater than the dynamic threshold value by passing an I/O request of the data element to the cache.

21. (Currently Amended) A data prefetch scheduling system comprising:

An LRU cache configured to communicate with a host; and

a remote prefetch module configured to communicate with the host and the LRU cache, intercept a stream of Input/Output (I/O) information from the host to the LRU cache to locate a requested data element and further configured to determine whether to schedule a prefetch of data into the LRU cache, wherein the determination to prefetch data is at least partially determined based on whether a data element preceding a requested data element is present in the LRU cache;

a modeling module operating within the remote prefetch module configured to model the LRU cache, including providing a model of data elements currently stored within the LRU cache, wherein each data element is assigned a priority value according to its history; and

a prefetch request module configured to request a data I/O from the LRU cache when the remote prefetch module determines that a prefetch is to be conducted.

22. (Original) The data prefetch scheduling system of claim 21, wherein the cache comprises a least recently used (LRU) cache.

23. (Original) The data prefetch scheduling system of claim 22, wherein the LRU cache is a native LRU-only cache that is not internally modified.

24. (Previously Presented) The data prefetch scheduling system of claim 21, wherein the remote prefetch module further comprises a calculation module configured to compare the priority value assigned to a data element to a threshold value and determine whether to schedule a prefetch of the data element.

25. (Original) The data prefetch scheduling system of claim 21, wherein the remote prefetch module further comprises a dynamic threshold optimization configured to calculate and update a dynamic threshold used in determining whether to prefetch data.

26. (Original) The data prefetch scheduling system of claim 21, wherein the remote prefetch module is configured to model the cache for use in determining when to prefetch I/O data into the cache.

27. (Original) The data prefetch scheduling system of claim 21, wherein the remote prefetch module is configured to prefetch data into the cache according to a priority scheme that takes into account the run length of each sequential I/O stream.

28. (Cancelled).

29. (Currently Amended) A remote prefetch module for determining whether to schedule a prefetch of data into a cache of a computer system, the prefetch module comprising:

a modeling module configured to model dynamic operation of the cache;

wherein the modeling module is further configured to provide a model of data elements currently stored within the cache, wherein each data element is assigned a priority value according to its history and requested data elements are assigned a priority value based at least partially on whether a preceding data element is present in the cache; and

a calculation module configured to make a cache management decision based upon the model, wherein the cache management decision is at least partially determined based on intercepting a stream of I/O information from a host to the cache to locate a requested data element and is further based on whether a data element preceding a the requested data element is present in the cache.

30. (Currently Amended) A computer station on a computer network, wherein the computer station is configured to communicate with a cache coupled to a storage device of the computer network, the computer station comprising:

a processor; and

a memory configured to store data structures comprising:

a modeling module configured to model dynamic operation of the cache;

wherein the modeling module is further configured to provide a model of data elements currently stored within the cache, wherein each data element is assigned a priority value according to its history and requested data elements are assigned a priority value based at least partially on whether a preceding data element is present in the cache; and

a calculation module configured to make a cache management decision based upon the model, wherein the cache management decision is at least partially determined based on intercepting a stream of I/O information from a host to the cache to locate a requested data element and is further based on whether a data element preceding a requested data element is present in the cache.

31. (Currently Amended) A computer readable medium comprising executable data structures configured to carry out a method for scheduling prefetches into a cache of a data storage system, the method comprising:

remotely modeling dynamic operation of the cache in a model, the model including a model of data elements currently stored within the cache;

assigning a priority value to modeled data elements according to their history;

assigning a priority value to a requested data element based at least partially on whether a preceding data element is present in the cache;

making a cache management decision based upon the model ~~and~~;

executing prefetches into the cache in response to select cache management decisions;

making the requested data element a youngest member of the cache.

32. (Currently Amended) A data prefetch scheduling system comprising:

a means for remotely modeling dynamic operation of ~~the~~ an LRU cache in a model ; , the remotely modeling including providing a model of data elements currently stored within the LRU cache;

a means for assigning a priority value to modeled data elements according to their history;

a means for assigning a priority value to a requested data element based at least partially on whether a preceding data element is present in the LRU cache;

a means for making a cache management decision based upon the model comprising;

a means for intercepting a request for a data element from a stream of Input/Output (I/O) data requests sent from a host and addressed to the LRU cache;

a means for determining whether to schedule a prefetch of a data element logically successive to the requested data element in accordance with contents of the LRU cache as indicated by the remote model; and

a means for executing prefetches into the LRU cache in response to select cache management decisions.